Byte-based Language Identification with Deep Convolutional Networks

Johannes Bjerva
University of Groningen
The Netherlands
j.bjerva@rug.nl

arXiv:1609.09004v1 [cs.CL] 28 Sep 2016

Abstract

We report on our system for the shared task on discriminating between similar languages (DSL 2016). The system uses only byte representations in a deep residual network (ResNet). The system, named ResIdent, is trained only on the data released with the task (closed training). We obtain 84.88% accuracy on subtask A, 68.80% accuracy on subtask B1, and 69.80% accuracy on subtask B2. A large difference in accuracy on development data can be observed with relatively minor changes in our network's architecture and hyperparameters. We therefore expect fine-tuning of these parameters to yield higher accuracies.

1 Introduction

Language identification is an unsolved problem, certainly in the context of discriminating between very similar languages. This problem is tackled in the Discriminating between Similar Languages (DSL) series of shared tasks (Zampieri et al., 2014; Zampieri et al., 2015). Most successful approaches to the DSL shared task in previous years have relied on settings containing ensembles of classifiers (Goutte et al., 2016). These classifiers often use various combinations of features, mostly based on word and/or character n-grams.

We are interested in exploring a single methodological aspect in the current edition of this shared task (Malmasi et al., 2016). We aim to investigate whether reasonable results for this task could be obtained by applying recently emerged neural network architectures, coupled with sub-token input representations. To address this question, we will look at convolutional neural networks (CNNs) and recurrent neural networks (RNNs). A recent development is the emergence of deep residual networks (ResNets), a building block for CNNs which have been proven to yield excellent results in image classification tasks (He et al., 2015; He et al., 2016). ResNets consist of several stacked residual units, which can be thought of as a collection of convolutional layers coupled with a 'shortcut' which aids the propagation of the signal in a neural network. This allows for the construction of much deeper networks, since keeping a 'clean' information path in the network facilitates optimisation (He et al., 2016). ResNets have recently shown state-of-the-art performance for Semantic and Part-of-Speech tagging (Bjerva et al., 2016), and have also seen other recent use in NLP (Östling, 2016; Conneau et al., 2016). However, no previous work has attempted to apply ResNets to language identification tasks. Additionally, most successful systems in previous editions of the DSL shared task have also applied word representations, whereas we focus solely on byte-level representations.

2 Method

Several previous approaches in the DSL shared tasks have formulated the task as a two-step classification, first identifying the language group, and then the specific language (Zampieri et al., 2015). Instead of taking this approach, we formulate the task as a multi-class classification problem, with each language / dialect representing a separate class. Our system is a deep neural network consisting of a bidirectional Gated Recurrent Unit (GRU) network at the upper level, and a Deep Residual Network (ResNet) at the lower level (Figure 1). The inputs of our system are byte-level representations of each input sentence,

with byte embeddings which are learnt during training. This architecture was based on the sequence-to-sequence labelling architecture used in Bjerva et al. (2016), modified for the task of language identification. Our system is implemented in Keras using the Tensorflow backend (Chollet, 2015; Abadi et al., 2016).

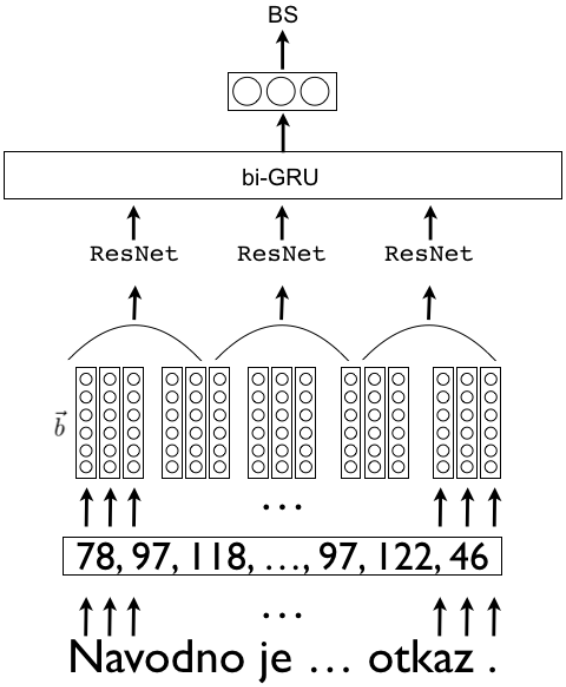

Figure 1: Model architecture: ResNet with byte representations ($\vec{b}$), with a bi-GRU at the upper level. The input example sequence is converted to a sequence of byte embeddings, which is treated by the ResNet, followed by the bi-GRU, finally yielding the language id *BS* (Bosnian).

2.1 Gated Recurrent Unit Networks

GRUs (Cho et al., 2014) are a recently introduced variant of RNNs, and are designed to prevent vanishing gradients, thus being able to cope with longer input sequences than vanilla RNNs. GRUs are similar to the more commonly-used Long Short-Term Memory networks (LSTMs), both in purpose and implementation (Chung et al., 2014). A bi-directional GRU makes both forward and backward passes over sequences, and can therefore use both preceding and succeeding contexts to predict a tag (Graves and Schmidhuber, 2005; Goldberg, 2015). Bi-directional GRUs and LSTMs have been shown to yield high performance on several NLP tasks, such as POS and semantic tagging, named entity tagging, and chunking (Wang et al., 2015; Yang et al., 2016; Plank et al., 2016; Bjerva et al., 2016).

2.2 Deep Residual Networks

Deep Residual Networks (ResNets) are built up by stacking residual units. A residual unit can be expressed as:

$$\begin{aligned} y_l &= h(x_l) + \mathcal{F}(x_l, \mathcal{W}_l), \\ x_{l+1} &= f(y_l), \end{aligned} \tag{1}$$

where x_l and x_{l+1} are the input and output of the l-th layer, $\mathcal{W}_l$ is the weights for the l-th layer, and $\mathcal{F}$ is a residual function (He et al., 2016), e.g., the identity function (He et al., 2015), which we also use in our experiments. ResNets can be intuitively understood by thinking of residual functions as paths through which information can propagate easily. This means that, in every layer, a ResNet learns more complex feature combinations, which it combines with the shallower representation from the previous layer. This architecture allows for the construction of much deeper networks. ResNets have recently been found to yield impressive performance in image recognition tasks, with networks as deep as 1001 layers (He et

al., 2015; He et al., 2016), and are thus an interesting and effective alternative to simply stacking layers. In this paper we use the *assymetric* variant of ResNets as described in Equation 9 in He et al. (2016):

$$x_{l+1} = x_l + \mathcal{F}(\hat{f}(x_l), \mathcal{W}_l). \tag{2}$$

2.3 System Description

We represent each sentence using a byte-based representation (S_b). This representation is a 2-dimensional matrix $S_b \in \mathbb{R}^{s \times d_b}$, where s is the zero-padded sentence length and d_b is the dimensionality of the byte embeddings. Byte embeddings are passed through a ResNet prior to being treated by the bi-GRU.

2.3.1 Hyperparameters

We use rectified linear units (ReLUs) for all activation functions (Nair and Hinton, 2010), and apply dropout (DO) with $p = 0.1$ to both input weights and recurrent weights in the bi-GRU (Srivastava et al., 2014). Within the ResNet, we apply batch normalisation (BN) followed by DO with $p = 0.5$ after each layer (Ioffe and Szegedy, 2015). One residual block in our architecture is defined by: a $1d$ convolution of length 8, followed by BN, a ReLu and DO which is followed by another $1d$ convolution of length 4, followed by BN, a ReLu and DO. The output of this residual block is concatenated to its own input, which is then pooled (factor 2) and used as an input for the next residual block. We set both d_b and d_w to 64. All GRU layers have 100 hidden units. All experiments were run with early stopping monitoring validation set loss, using a maximum of 50 epochs. We use a batch size of 100. Optimisation is done using the ADAM algorithm (Kingma and Ba, 2014), with the categorical cross-entropy loss function as training objective.

For the B tasks, we train the model in the same way as for the A tasks. Only a handful of instances ($n \approx 5$) per B run are classified as belonging to a language which is not in the B group. These cases are automatically converted to being in the class *hr*. For the B tasks, we also perform a simple clean-up of the data. We first remove all hyperlinks, hashtags and usernames from the text with a simple regex-based script. We then remove all tweets classified as English.

3 Results

Table 1: Results for all runs in subtasks A, B1 and B2 (closed training).

Test Set	Track	Run	Accuracy	F1 (micro)	F1 (macro)	F1 (weighted)
A	closed	Baseline	0.083			
A	closed	run1	0.8462	0.8462	0.8415	0.8415
A	closed	run2	0.8324	0.8324	0.8272	0.8272
A	closed	run3	**0.8488**	0.8488	0.8467	0.8467
B1	closed	Baseline	0.020			
B1	closed	run1	0.682	0.682	0.6802	0.6802
B1	closed	run2	0.676	0.676	0.6708	0.6708
B1	closed	run3	**0.688**	0.688	0.6868	0.6868
B2	closed	Baseline	0.020			
B2	closed	run1	0.684	0.684	0.6788	0.6788
B2	closed	run2	**0.698**	0.698	0.6942	0.6942
B2	closed	run3	0.664	0.664	0.6524	0.6524

The system used for runs 1, 2 and 3 contain five, four and three residual blocks respectively. For subtasks A and B1, run 3 results in the best accuracy on test, whereas run 2 results in the best accuracy on B2. The results are shown in Table 1. Figure 2 contains a confusion matrix for the results in subtask A.

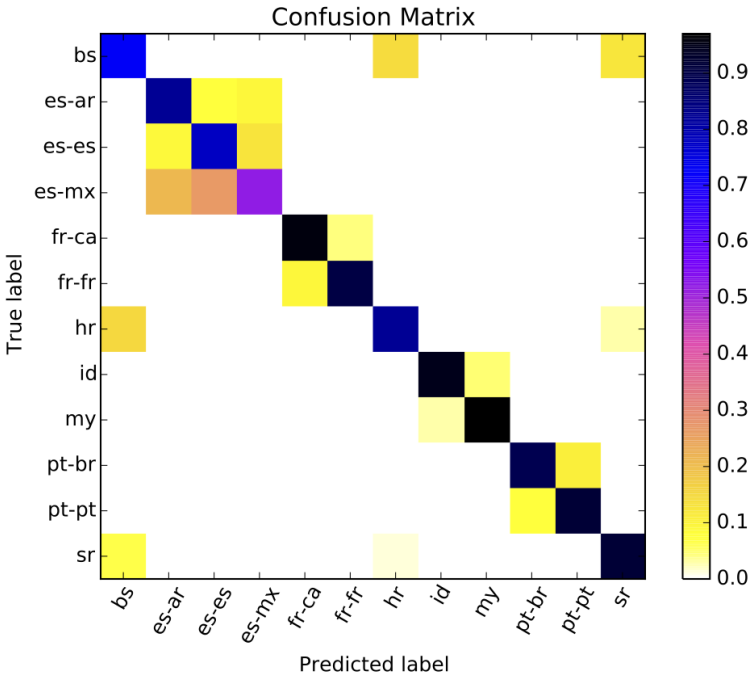

Figure 2: Confusion matrix, closed run 3, on test set A.

4 Discussion

Figure 2 contains a confusion matrix of our best run for subtask A. Judging from the confusion matrix, our system has very low confusion between language groups. However, confusion can be observed within all groups. Although the system achieves reasonable performance, there is a large gap between our system and the best performing systems. This can to some extent be explained by limitations caused by our implementation.

The largest limiting factor can be found in the fact that we only allowed our system to use the first ca. 384 bytes of each training/testing instance. For the training and development set, and subtask A, this was no major limitation, as this allowed us to use more than 90% of the available data. However, for subtasks B1 and B2, this may have seriously affected the system's performance. Additionally, we restricted our system to using only byte embeddings as input. Adding word-level representations into the mix, would likely increase system performance.

We also observed considerable differences in development accuracy when changing hyperparameters of our network in relatively minor ways. For instance, altering the patch sizes used in our CNNs had a noticeable impact on validation loss. However, altering the amount of residual blocks used, did not have a large effect on results. The neural network architecture, as well as most of the hyperparameters, were tuned on an altogether different task (semantic tagging), and adapted for the current task. Further fine tuning of the network architecture and hyperparameters for this task would therefore likely lead to narrowing the performance gap.

5 Conclusions

We implemented a language identification system using deep residual networks (ResNets) coupled with a bidirectional Gated Recurrent Unit network (bi-GRU), using only byte-level input representations. In the DSL 2016 shared task, we achieved a reasonable performance, with 84.88% accuracy on subtask A, 68.80% accuracy on subtask B1, and 69.80% accuracy on subtask B2. Although reasonable performance was achieved, further fine tuning of input representations and system architecture would likely improve performance.

References

[Abadi et al.2016] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Gregory S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian J. Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Józefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mane, Rajat Monga, Sherry Moore, Derek Gordon Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul A. Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda B. Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.

[Bjerva et al.2016] Johannes Bjerva, Barbara Plank, and Johan Bos. 2016. Semantic Tagging with Deep Residual Networks. *arXiv preprint arXiv:1609.07053*.

[Cho et al.2014] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*.

[Chollet2015] François Chollet. 2015. Keras. `https://github.com/fchollet/keras`.

[Chung et al.2014] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.

[Conneau et al.2016] Alexis Conneau, Holger Schwenk, Loïc Barrault, and Yann Lecun. 2016. Very deep convolutional networks for natural language processing. *arXiv preprint arXiv:1606.01781*.

[Goldberg2015] Yoav Goldberg. 2015. A primer on neural network models for natural language processing. *arXiv preprint arXiv:1510.00726*.

[Goutte et al.2016] Cyril Goutte, Serge Léger, Shervin Malmasi, and Marcos Zampieri. 2016. Discriminating Similar Languages: Evaluations and Explorations.

[Graves and Schmidhuber2005] Alex Graves and Jürgen Schmidhuber. 2005. Framewise phoneme classification with bidirectional lstm and other neural network architectures. *Neural Networks*, 18(5):602–610.

[He et al.2015] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*.

[He et al.2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Identity mappings in deep residual networks. *arXiv preprint arXiv:1603.05027*.

[Ioffe and Szegedy2015] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.

[Kingma and Ba2014] Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

[Malmasi et al.2016] Shervin Malmasi, Marcos Zampieri, Nikola Ljubešić, Preslav Nakov, Ahmed Ali, and Jörg Tiedemann. 2016. Discriminating between similar languages and arabic dialect identification: A report on the third dsl shared task. In *Proceedings of the 3rd Workshop on Language Technology for Closely Related Languages, Varieties and Dialects (VarDial)*, Osaka, Japan.

[Nair and Hinton2010] Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 807–814.

[Östling2016] Robert Östling. 2016. Morphological reinflection with convolutional neural networks. In *Proceedings of the 2016 Meeting of SIGMORPHON*, Berlin, Germany. Association for Computational Linguistics.

[Plank et al.2016] Barbara Plank, Anders Søgaard, and Yoav Goldberg. 2016. Multilingual part-of-speech tagging with bidirectional long short-term memory models and auxiliary loss. In *Proceedings of ACL 2016, arXiv preprint arXiv:1604.05529*.

[Srivastava et al.2014] Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.

[Wang et al.2015] Peilu Wang, Yao Qian, Frank K Soong, Lei He, and Hai Zhao. 2015. A unified tagging solution: Bidirectional lstm recurrent neural network with word embedding. *arXiv preprint arXiv:1511.00215*.

[Yang et al.2016] Zhilin Yang, Ruslan Salakhutdinov, and William Cohen. 2016. Multi-task cross-lingual sequence tagging from scratch. *arXiv preprint arXiv:1603.06270*.

[Zampieri et al.2014] Marcos Zampieri, Liling Tan, Nikola Ljubešić, and Jörg Tiedemann. 2014. A report on the dsl shared task 2014. In *Proceedings of the First Workshop on Applying NLP Tools to Similar Languages, Varieties and Dialects (VarDial)*, pages 58–67, Dublin, Ireland.

[Zampieri et al.2015] Marcos Zampieri, Liling Tan, Nikola Ljubešić, Jörg Tiedemann, and Preslav Nakov. 2015. Overview of the dsl shared task 2015. In *Proceedings of the Joint Workshop on Language Technology for Closely Related Languages, Varieties and Dialects (LT4VarDial)*, pages 1–9, Hissar, Bulgaria.